



Special issue on learning for human–robot collaboration

Leonel Rozo¹ · Heni Ben Amor² · Sylvain Calinon³ · Anca Dragan⁴ · Dongheui Lee^{5,6}

Published online: 26 April 2018

© Springer Science+Business Media, LLC, part of Springer Nature 2018

Once isolated behind safety fences, robots are now making their way into spaces shared with people: not only on manufacturing production lines, but also our houses, museums, or hospitals. In these spaces, collaboration between humans and robots becomes vital, but also raises new challenges for robotics algorithms. A collaborative robot must be able to assist humans in a large diversity of tasks, understand its collaborator's intentions as well as communicate its own, predict human actions to adapt its behavior accordingly, and decide when it can lead the task or instead follow the user. All these aspects require the robot to adapt. It needs to execute different tasks, and rapidly adapt to the user's actions and requirements. This adaptation requirement makes learning a crucial feature for collaborative robots.

The goal of this special issue is to document and highlight recent progress in robot learning for human–robot collaboration (HRC), covering a diversity of articles that reflect the state-of-the-art in the field. Following an open call for papers,

we received more than fifty submissions, among which ten articles were selected for the special issue, covering exciting applications spanning co-manipulation, medical robotics, and social behaviors. These works and five additional papers compose a fascinating online topical collection that reports a larger diversity of research advancements in learning for HRC.

In **Progress and Prospects of the Physical Human–Robot Collaboration**, Ajoudani, Zanchettin, Ivaldi, Albu-Schäffer, Kosuge and Khatib review the state-of-the-art on intermediate human–robot bi-directional interfaces, robot control modalities, systems stability, benchmarking, and relevant use cases. The article extends views on the required future developments in the realm of physical human–robot collaboration. The authors provide a thorough overview of the pioneering methodologies that aim at achieving intuitive and seamless HRC, potentially relying on a minimum degree of task-related pre-programming, where the relevance of robot learning stands out. They also provide a list of potential applications and relevant use cases ranging from domestic to industrial environments.

Turn-taking prediction, a synchronization aspect in HRC, is the capability to comprehend the on-going task progress and predict where, when and how to seize the next turn during multi-agent collaborations. In **Early Prediction for Physical Human–Robot Collaboration in the Operating Room**, Zhou and Wachs show the importance of fluent and natural turn-taking regulations in team performance, which lead to better social connections among team members. The requirement for early turn-taking prediction stands out more clearly in high-risk and high-paced tasks like surgery, where the scrub nurse and the surgeon perform fast, accurate and highly coordinated turn-taking actions when exchanging surgical instruments. In this context, the authors tackle the problem of designing a fully functional robotic scrub nurse, and propose a computational framework for early turn-taking prediction built on Long Short-Term Memory networks and Dempster–Shafer theory for uncertain sensor fusion. The approach is evaluated on a simulated surgical procedure dataset, and is

✉ Leonel Rozo
leonel.rozo@iit.it

Heni Ben Amor
hbenamor@asu.edu

Sylvain Calinon
sylvain.calinon@idiap.ch

Anca Dragan
anca@berkeley.edu

Dongheui Lee
dhlee@tum.de

¹ Istituto Italiano di Tecnologia, Via Morego 30, 16163 Genoa, Italy

² Arizona State University, 699 S Mill Ave, Tempe, AZ 85281, USA

³ Idiap Research Institute, Rue Marconi 19, 1920 Martigny, Switzerland

⁴ University of California, Berkeley, Berkeley, CA 94720-1758, USA

⁵ Technical University of Munich, Karlstraße 45, 80333 Munich, Germany

⁶ German Aerospace Center (DLR), Münchenerstraße 20, 82234 Weßling, Germany

found to be superior to its algorithmic counterparts, and better than the human baseline when little partial input is given.

Shared autonomy is on its own a relevant form of collaboration, and medical robotics has also inspired the research on shared autonomy reported in the paper **Skill-Based Human–Robot Cooperation in Teleoperated Path Tracking** by Enayati, Ferrigno and De Momi. The paper argues that a robot’s assistance during shared autonomy should depend on the operator’s skill level. The robot should aid novices, but not unnecessarily restrict experts. This work introduces a skill estimation system for path tracking tasks, and puts the idea to the test in a user study. The results support that customizing assistance based on operator’s skill improves the shared autonomy experience.

Human–robot collaborative manipulation entails bi-directional force exchange, mutual adaptation, variable compliance, among others. Such scenario served as an inspiration for several works in this special issue. For example, in **Robot Adaptation to Human Physical Fatigue in Human–Robot Co-Manipulation**, Peternel, Tsagarakis, Caldwell and Ajoudani propose to estimate human fatigue, and to adjust how much a robot is helping in a human–robot collaborative manipulation task based on how tired the person gets. The robot begins by acting as a follower, and observes the trajectory, stiffness, and/or forces that the user prefers. As the user is performing the task, the robot collects these examples and also monitors the estimated human shoulder muscle activation via electromyography (EMG). Once the user’s fatigue reaches a predefined threshold, the robot takes control of the task and follows the learned stiffness/force and trajectories.

In **Human–Robot Cooperation with Compliance Adaptation along the Motion Trajectory**, Nemeč, Likar, Gams and Ude introduce a learning from demonstration interface for kinesthetic teaching of compliant co-manipulation tasks. Since it is difficult to demonstrate precise motions at high speed, the paper proposes to split the teaching procedure into two parts. First, the human demonstrates the desired trajectory, and the robot infers the desired stiffness from the variance in the demonstrations. Second, the human demonstrates the desired speed profile, with trajectory guidance provided by the robot based on the first training step. The proposed scheme was experimentally verified in a human–robot collaborative transportation task, which may be further exploited in assembly processes in production plants or in civil engineering for transportation of heavy and bulky objects.

In **Co-manipulation with a Library of Virtual Guiding Fixtures**, Raiola, Sanchez Restrepo, Chevalier, Rodriguez-Ayerbe, Lamy, Tliba and Stulp address the problem of learning virtual guiding fixtures in HRC. These guides are built on a probabilistic representation based on Gaussian mixture regression, and are aimed at constraining the movements

of a robot to task-relevant trajectories defined by the user. By using the analogy of a ruler used as a physical guide to draw lines, they emphasize that robots can be used to define virtual guides of complex shapes, enabling safer co-manipulation for industrial tasks. Whereas previous work mostly considered guiding fixtures for single tasks, this paper addresses the problem of creating a library of guiding fixtures for multiple tasks, by selecting the guides online and incrementally refining them. The approach is demonstrated in a user study with an industrial pick-and-place task, showing that a library of guiding fixtures can provide an intuitive haptic interface for joint human–robot completion of tasks, improving task execution time, and reducing mental workload and errors.

Handovers are complex interactions where agents coordinate in time and space to transfer control of an object. This challenging interaction is the focus of the paper **One-Shot Learning of Human–Robot Handovers with Triadic Interaction Meshes** by Vogt, Stepputtis, Jung and Ben Amor, where it is proposed to learn human–robot handover interaction from observing one demonstration of human–human handover. The relevant information, such as joint correlation and spatial relationships of two humans and a handing-over object, is extracted from a single demonstration. Triadic interaction meshes are constructed to model interaction between two agents and an object, and are subsequently used for robot’s adaptive motion generation in human–robot handover scenarios.

Human–Robot social interactions also play an essential role in extending the use of robots in daily life, where robot learning can be exploited to master different types of behaviors. In this context, the paper **Learning Proactive Behavior for Interactive Social Robots** by Liu, Glas, Kanda and Ishiguro, addresses the learning of both reactive and proactive robot behaviors from human–human interactions. The authors first introduce the concept of a “yield action” that enables the robot to identify opportunities for a proactive action to be generated. Since proactive behaviors are often sensitive to the context of the interaction, interaction history is incorporated into the inputs of a deep neural network, along with variables describing natural language and the state of the interactive agents. An attention mechanism is also used in the learning system, which has the ability to “attend” and learn which parts of the interaction history are relevant. Offline analysis and live interactions with users in a camera shop scenario showed that the proposed system can effectively reproduce proactive behaviors that were perceived as socially-appropriate by the participants.


The paper **Hierarchical Emotional Episodic Memory for Social Human–Robot Collaboration** by Lee and Kim investigates how human emotional states can be represented and anticipated in order to improve the HRC experience. A critical component to achieve these goals is a computational model of episodic memory based on Adaptive Resonance

Theory. The resulting framework allows robots to learn correlations between the user's emotional state and the executed collaborative actions. In turn, the introduced framework and algorithms can be used in a variety of tasks, in particular in domestic environments, in order to endow robots with social capabilities.

Finally, a fresh look at interaction learning methods for collaborative tasks is given in the paper **Efficient Behavior Learning in Human–Robot Collaboration** by Munzer, Toussaint and Lopes. In particular, the paper discusses a relational approach to learning task requirements and user preferences. A subset of first-order logic is used to explicitly represent knowledge about the collaborative task. At the same time, a reinforcement learning method is used to iteratively update the representation, so as to increase efficiency during the joint task. The approach aims at bridging the divide between symbolic and numerical approaches to artificial intelligence for robot systems.

Collectively, these ten papers illustrate a broad range of human–robot collaboration applications. They provide a compilation of the large variety of challenges that are currently being investigated, where robot learning emerges as a promising approach to tackle this diversity.

Acknowledgements We would like to thank the reviewers who provided insightful comments on the papers submitted to this special issue. We would also like to thank Dr. Gaurav Sukhatme, the Editor-in-Chief of *Autonomous Robots*, for his assistance with this special issue. Thanks also go to Sarvagnan Subramanian for helping us with the process and maintaining interactions with the reviewers and the authors. Last but not least, we would like to thank all authors for contributing their work to this special edition.



Leonel Rozo is currently the Team Leader of the Learning and Interaction Group at the Department of Advanced Robotics (ADVR), Istituto Italiano di Tecnologia. He was a Postdoctoral Researcher at the same institution from 2013 to 2016. He received his Ph.D. in Robotics and Machine Learning (2013) from the Polytechnic University of Catalonia (Barcelona, Spain). His research interests cover robot learning, human–robot interaction, machine learning and optimal control for robotics. Link:

<http://leonelrozo.weebly.com/>.



Heni Ben Amor is an Assistant Professor of Robotics at Arizona State University and the Director of the ASU Interactive Robotics Laboratory. He was a research scientist at Georgia Tech, a Postdoctoral Researcher at the Technical University Darmstadt (Germany), and a visiting research scientist in the Intelligent Robotics Lab at the University of Osaka (Japan). He has won the Daimler-Benz Award, the NSF CAREER award, and various best paper awards at international conferences.

His primary research interests lie in the fields of artificial intelligence, machine learning, robotics, and human–robot interaction. Link: <http://henibenamor.weebly.com/>.



Sylvain Calinon is a Senior Researcher at the Idiap Research Institute, a Lecturer at the Ecole Polytechnique Fédérale de Lausanne (EPFL), and an External Collaborator at the Department of Advanced Robotics (ADVR), Italian Institute of Technology (IIT). He was previously a Team Leader at ADVR, IIT and a Postdoc at the Learning Algorithms and Systems Laboratory, EPFL, where he obtained his Ph.D. in 2007. His research interests cover robot learning and human–robot interaction,

with recognition including Best Paper Awards in the *Journal of Intelligent Service Robotics* (2017) and at IEEE Ro-Man'2007, as well as Award Finalists at ICRA'2016, ICIRA'2015, IROS'2013 and Humanoids'2009. Link: <http://calinon.ch>.



Anca Dragan is an Assistant Professor in the Electrical Engineering and Computer Sciences Department at UC Berkeley. She runs the InterACT Lab, which focuses on algorithms that enable robots to work with, around, and in support of people. Anca got her Ph.D. from the Robotics Institute at Carnegie Mellon (2015). She helped found and serves on the steering committee for the Berkeley AI Research Lab, and is a co-Pi on the Center for Human-Compatible AI. She has won an

NSF CAREER award, the Okawa Foundation award, is a Sloan Fellow, and was on MIT Tech Review's 35 innovators under 35 list. Her research has been featured in *The Atlantic*, *New Scientist*, *WIRED*, *IEEE Spectrum*, and *NPR*.



Dongheui Lee is an Associate Professor at the Department of Electrical and Computer Engineering, Technical University of Munich (TUM) and a Head of Human-Centered Assistive Robotics Group at the Institute of Robotics and Mechatronics at DLR, Germany. She received her Ph.D. degree (2007) from the University of Tokyo (UTokyo), Japan, and worked as a Research Scientist at Korea Institute of Science and Technology from 2001 to 2004. In 2007 she joined the Center of

Information and Robot Technology at UTokyo as a Project Assistant Professor. In 2009 she joined TUM as an Assistant Professor. She was awarded for a Carl-von-Linde Fellowship at TUM (2011) and a Helmholtz Professorship Prize (2015). Her research interests include human motion understanding, human–robot interaction, machine learning in robotics, and assistive robotics.